

H04B1/707

VN. 8272

A Low-Power DSP Core-Based Software Radio Architecture

James E. Gunn, Senior Member, IEEE, Kenneth S. Barron, and William Ruczczyk

P.574 - 590 = (17) PD, 00/04/99

XP-000824304

Abstract— This paper describes an approach to developing a low-power digital signal processor (DSP) subsystem architecture for advanced software radio platforms. The architecture is intended to support next-generation wide-band spread-spectrum military waveforms. The methodology illustrates how a next-generation programmable DSP core forms the basis for an application-specific integrated circuit (ASIC). It also shows how semiconductor technologies can be integrated into such chips to achieve algorithm performance while minimizing subsystem power consumption. The ASIC is run-time configurable to maintain high flexibility. The range of RF channel modulation ("waveforms") and air interfaces is intended to include both wide-band and traditional narrow-band waveforms. Estimated gate counts and power-consumption estimates are presented. DSP circuit-design and power-management strategies necessary to achieve low-power operation are presented. While the architecture discussion focuses on military waveforms, the approach is also applicable to commercial waveforms.

Index Terms— Communication, configurable ASIC, DSP architecture, low power, programmable DSP, software radio, wireless.

I. INTRODUCTION

SOFTWARE radios migrate the traditional hard-wired radio platforms to flexible software radio platforms that can support multiple modulation waveforms and multiple air interface standards. This approach should allow the graceful evolution of the technology over time. Thus, software radio hardware platforms can serve a range of applications including: analog cellular; digital cellular/personal communications services (PCS); advanced wide-band, spread-spectrum military waveforms; legacy narrow-band military waveforms; navigation waveforms (e.g., the global positioning system); and emergency preparedness, public safety, and other waveforms. Depending on the waveform(s), architecture, and implementation, a single software radio platform could have the flexibility to support a broad range of such waveforms.

Three general classes of programmable digital and software radio architectures are emerging: base-station, mobile, and battery-powered handheld units. Typically, base-stations support large numbers of channels and users with only a few types

of waveforms (e.g., cellular). They have less severe power and space constraints than the other classes. Conversely, battery-powered handheld units typically support a single user, have aggressive power and space constraints, and typically support one or two services. Users of handheld units have traditionally carried multiple devices to access multiple services (e.g., cellular telephone and a pager). Users seem to want expanded, single-platform, multiservice flexibility, encompassing multiple services (e.g., cellular/PCS, paging, data networks, private dispatch, military networks, and perhaps limited video). Mobile units fall between these extremes, mounted on vehicles or easily transported. Mobile units typically address military and dispatch applications (e.g., police, taxi, fire, military vehicles, etc.).

This paper focuses on the DSP subsystem for emerging handheld software radio devices. This technology enables the support of multiple waveforms on a programmable, configurable hardware platform that consumes low power.

Section II summarizes communication waveforms and software radio concepts applicable to the design of the low-power DSP platform. It provides an overview of illustrative communication waveforms, emphasizing the characteristics that define DSP processing requirements. Section III provides a brief overview of relevant RF technologies, their function, and the related DSP subsystem integration concepts. Section IV gives an overview of data-acquisition requirements and the important technical parameters from the DSP subsystem perspective. Section V discusses low-power design and power-management techniques that support handheld, low-power implementations. In Section VI, a run-time configurable ASIC architecture is presented that processes throughput-intensive algorithms in support of a standard programmable DSP core. Section VII summarizes gate-count and power estimates. Section VIII summarizes the design methodology, and Section IX draws conclusions.

II. COMMUNICATION WAVEFORMS

The software radio uses RF circuits, DSP's, and microprocessors to implement wireless-communication-system functions, typically with the algorithmic structure of Fig. 1. Software radios have been canonically partitioned [1] into antenna, radio frequency (RF), intermediate frequency (IF) processing, baseband (BB), bitstream, and source segments. As RF, data acquisition, and digital signal processing (DSP) technologies have advanced, the trend has been to move the analog-to-digital converter (ADC) closer to the antenna and to do

Manuscript received February 11, 1998; revised June 23, 1998 and July 9, 1998. This work was supported in part by the U.S. Army Communications and Electronics Command and the Defense Advanced Research Projects Agency under Contract DAAB07-97-D016.

J. E. Gunn is at Box 833157, Richardson, TX 75083-3157 USA.

K. S. Barron and W. Ruczczyk are with Raytheon Systems Co., Dallas, TX USA.

Publisher Item Identifier S 0733-8716(99)02971-6.

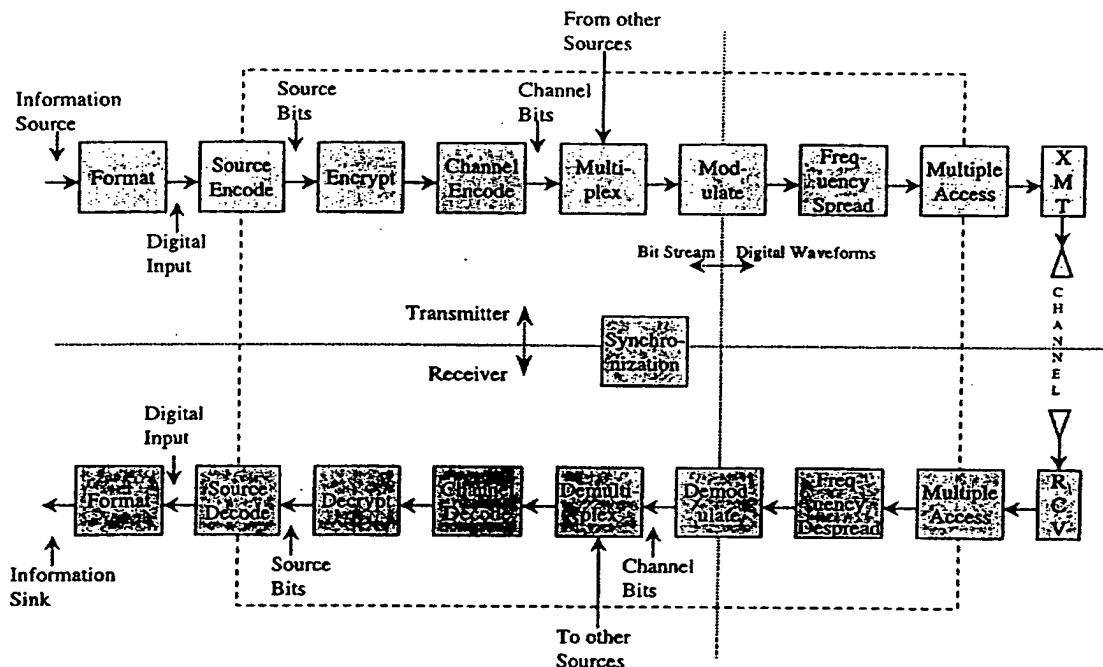


Fig. 1. Generic software radio algorithm block diagram (adapted from [2]).

TABLE I
REPRESENTATIVE DoD AND COMMERCIAL WAVEFORMS

WAVEFORM FAMILY	FREQUENCY RANGE	CHANNEL BANDWIDTH	MODULATION
FH HF	2 to 20 MHz	3 kHz	AM/FM
SINCGARS	30 to 88 MHz	25 kHz	FH (100 h/s)
Have Quick	225 to 400 MHz	25 kHz	FH
EPLRS	420 to 450 MHz	5 MHz	DSSS/FH
JTIDS	960 to 1310 MHz	6 MHz	DSSS/FH
GPS	1.5 GHz	10 MHz	DSSS
Speakeasy	All of the above		
SATCOM	7 GHz	variable	variable
IS-95	800 MHz, 1.9 GHz	1.25 MHz	DSSS/CDMA
IS-54/136	800 MHz	30 kHz	TDMA
ITU 3G	1.8 TO 2.1 GHz	5, 10, 15 MHz	DSSS/CDMA

more functions digitally. Therefore, the ability of a software radio architecture to support a communication waveform is predominantly determined by

- the largest instantaneous signal bandwidth (W);
- the frequency range and bandwidth of the RF;
- the ADC sampling rate (greater than $2W$);
- the maximum dynamic range;
- DSP throughput requirements including translation of IF to baseband, modulation, demodulation, coding, and decoding.

Table I lists the frequency band, bandwidth, and modulation format of some current DoD and commercial waveforms. Table II presents similar descriptions of waveforms

developed for Phase 1 of DARPA's small unit operations (SUO) project. These waveforms represent the state of the art in direct-sequence spread-spectrum (DSSS) and frequency-hopped spread-spectrum (FHSS) concepts. The architecture concepts presented in this paper address the DSP requirements implicit in these current and advanced waveforms. The emerging cellular/PCS wide-band, CDMA waveforms [3] currently being defined in international standards bodies should generally be addressed by this architecture as well.

The processing capacity required to modulate and code a waveform in the transmitter and to demodulate and decode the waveform in the receiver defines the DSP processing requirement of the modem. Fig. 2 graphically presents esti-

TABLE II
PROPOSED SMALL UNIT OPERATIONS WAVEFORMS [4]

Wave form	FREQUENCY RANGE	BAND-WIDTH	CHIP RATE	FH BAND-WIDTH	FH RATE	MODULATION FORMAT
1	30 to 88 MHz 225 to 400 MHz 0.8 to 2.45 GHz	Variable to 10 MHz	0.32 to 16 Mcps			Quasi Bandlimited MSK
2	77 to 88 MHz 225 to 400 MHz 1.8 to 2.0 GHz	Variable to 20 MHz	N/A	>= 100 MHz	100 hps	Wavelet (featureless, good side lobe suppression)
3	20 to 2000 MHz	Variable to 12 MHz	20 Mcps	200 MHz	400 hps	MSK
4	6 to 2000 MHz	Variable to 26 MHz	N/A	N/A	N/A	Transform Domain DSPN
5	10 to 2000 MHz	1.6 MHz	680 Kcps I & Q	Maximum 60 MHz or 20% Carrier	1,200 hps	non-LPI: MSK LPI: Filtered DS

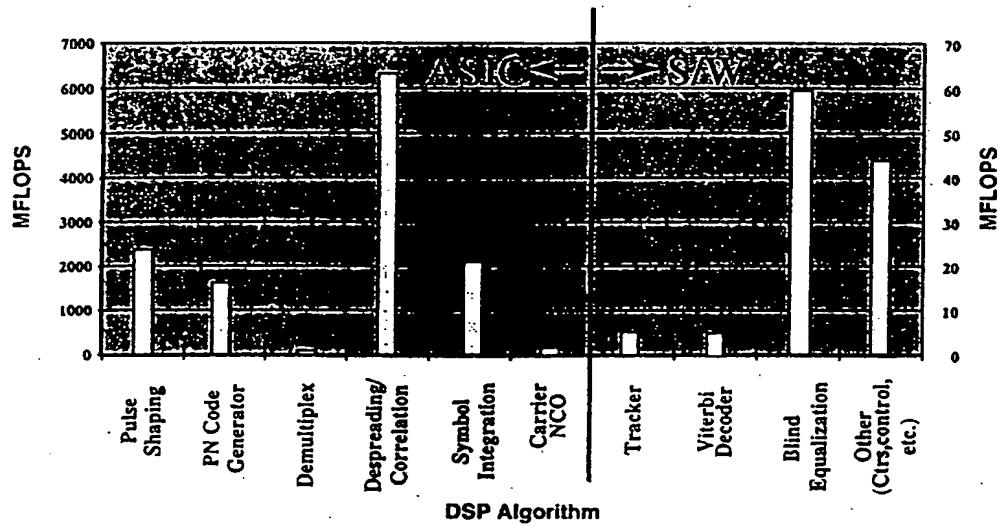


Fig. 2. Software radio receiver DSSS processing capacity estimates for 10 Mchips/s, 40 Msample/s (ADC), 9600 bits/s, 5 μ s delay spread; total: 12.78 GFLOPS, ASIC: 12.67 GFLOPS, S/W: 0.11 GFLOPS.

mated processing capacity for the key functions required to demodulate a proposed wide-band DSSS military waveform. This representative waveform has a DSSS chip rate of approximately 10 MCPS and a bandwidth of approximately 10–20 MHz (including frequency-domain sidelobes). The capacity requirements for spreading and despreading are not achievable on a programmable DSP. Thus, the DSP subsystem requires a communication preprocessor. This paper considers a design based on emerging DSP core and ASIC semiconductor technologies [5].

III. EMERGING RF SEGMENT ALTERNATIVES

Fig. 3 shows a typical system-level block diagram of a software radio transceiver. This superheterodyne receiver translates the RF to IF, digitizing the IF bandpass signal. This

architecture is well suited for traditional hardwired transceiver implementations that address both limited frequency range and fixed, typically narrow-band (e.g., 30 kHz cellular) channelization plans. Extended frequency range and variable channelization plans (narrow band and wide band) require multiple RF and IF filters and RF switches. Thus, flexibility is not easily implemented by using traditional design methods, and is not well suited for low-cost or low-power applications.

Fig. 4 shows three common RF receiver configurations. Fig. 4(a) shows the analog in-phase and quadrature (*I*, *Q*) superheterodyne receiver. The RF frequency is translated to baseband through one or more IF stages. The (*I*, *Q*) baseband signals are digitized and demodulated digitally in the DSP. The sampling rates of the ADC's need be at least the highest baseband signal frequency, whereas real sampling requires

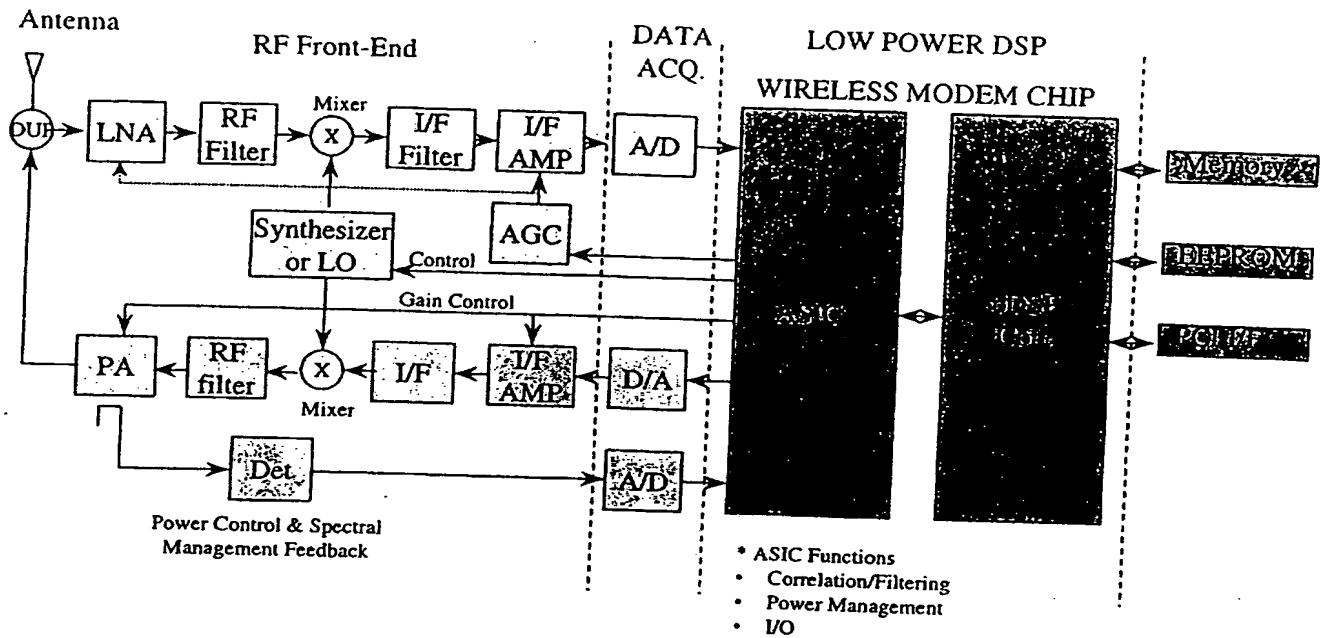


Fig. 3. Example system-level software radio transceiver block diagram [4], [6].

twice this sampling rate to meet the Nyquist criterion for unambiguous signal reconstruction. A second configuration is a passband superheterodyne receiver that translates in one or more IF stages to a final passband IF frequency where it is digitized, as shown in Fig. 4(b). The sampling rate of the passband signal must be at least twice the bandwidth of the upper cutoff frequency of the passband signal. Since the lower passband cutoff frequency will be greater than zero, typical sampling rates are two to three times the bandpass bandwidth. The sampling rate must be sufficiently greater than two times the signal passband bandwidth to ensure that sampled images of the positive and negative continuous spectrum [7] do not alias. Passband sampling is popular because only one ADC converter is required, simplifying the component configuration. For moderate bandwidths (e.g., 25 MHz), one can configure passband superheterodyne receivers using off-the-shelf ADC's (e.g., the Analog Devices 70 MHz ADC).

Fig. 4(c) shows a direct conversion (or homodyne or zero IF) receiver where the RF signal is translated directly to baseband. This requires that the numerically controlled oscillator (NCO) or synthesizer be locked to the carrier. The advantages of the direct conversion receiver include fewer signal translation steps, and the ability to use simpler analog filters cascaded with low-pass baseband digital filters in the DSP. This creates a more flexible (wider) tuning range and potentially greater channel bandwidths. The disadvantages include leakage from high-gain low-noise mixers, requirements for very high dynamic range analog components, the requirement for higher sensitivity than a comparable superheterodyne receiver, the need for precise I and Q phase balancing, dc offset cancellation, antenna isolation, and high-selectivity filters [8]. As a result, homodyne receivers are extremely challenging to implement.

Traditionally, tuning and translation to/from IF and baseband have been accomplished in the analog RF and IF seg-

ments [9], while (de)modulation and (de)coding have been done in the baseband DSP subsystem. With emerging commercial and military application for single platforms covering 2–3000 MHz having configurable modulation bandwidths, advanced RF conversion is more frequently based on passband superheterodyne or direct conversion receiver designs. Ongoing research on micro electromechanical systems (MEMS) [10], [11] offers significant miniaturization and power reduction potential for the filters, oscillators, and switches required for configurable RF IF segments that support the desired frequency bands and channel bandwidths.

The superheterodyne receiver with real, (I , Q), and bandpass sampling is applicable to low power software radios. Homodyne receivers are also applicable, with the advantage of reduced parts count, but at the risk of introducing artifacts into the baseband. A flexible ASIC-DSP core entails the flexibility to accommodate any of these RF approaches.

IV. SIGNAL CONVERSION

Signal conversion refers to analog-to-digital conversion (ADC) and its inverse, digital-to-analog conversion (DAC). The theory of sampling and quantization (i.e., digitization) of analog signals has a rich history [12]. From an implementation perspective, the ADC is more demanding than the DAC. In particular, ADC's have generally consumed considerable power, posing technical challenges for low-power software radios.

A technical discussion of ADC's for software radios is provided in a companion paper in this issue [13]. This includes an analysis of sample-and-hold and conversion circuits, aperture uncertainty, jitter, number of bits of resolution, precision, and related issues. The design issue of primary concern in this paper is power dissipation. Due to the need to rapidly change voltage levels in such a way that the power in the

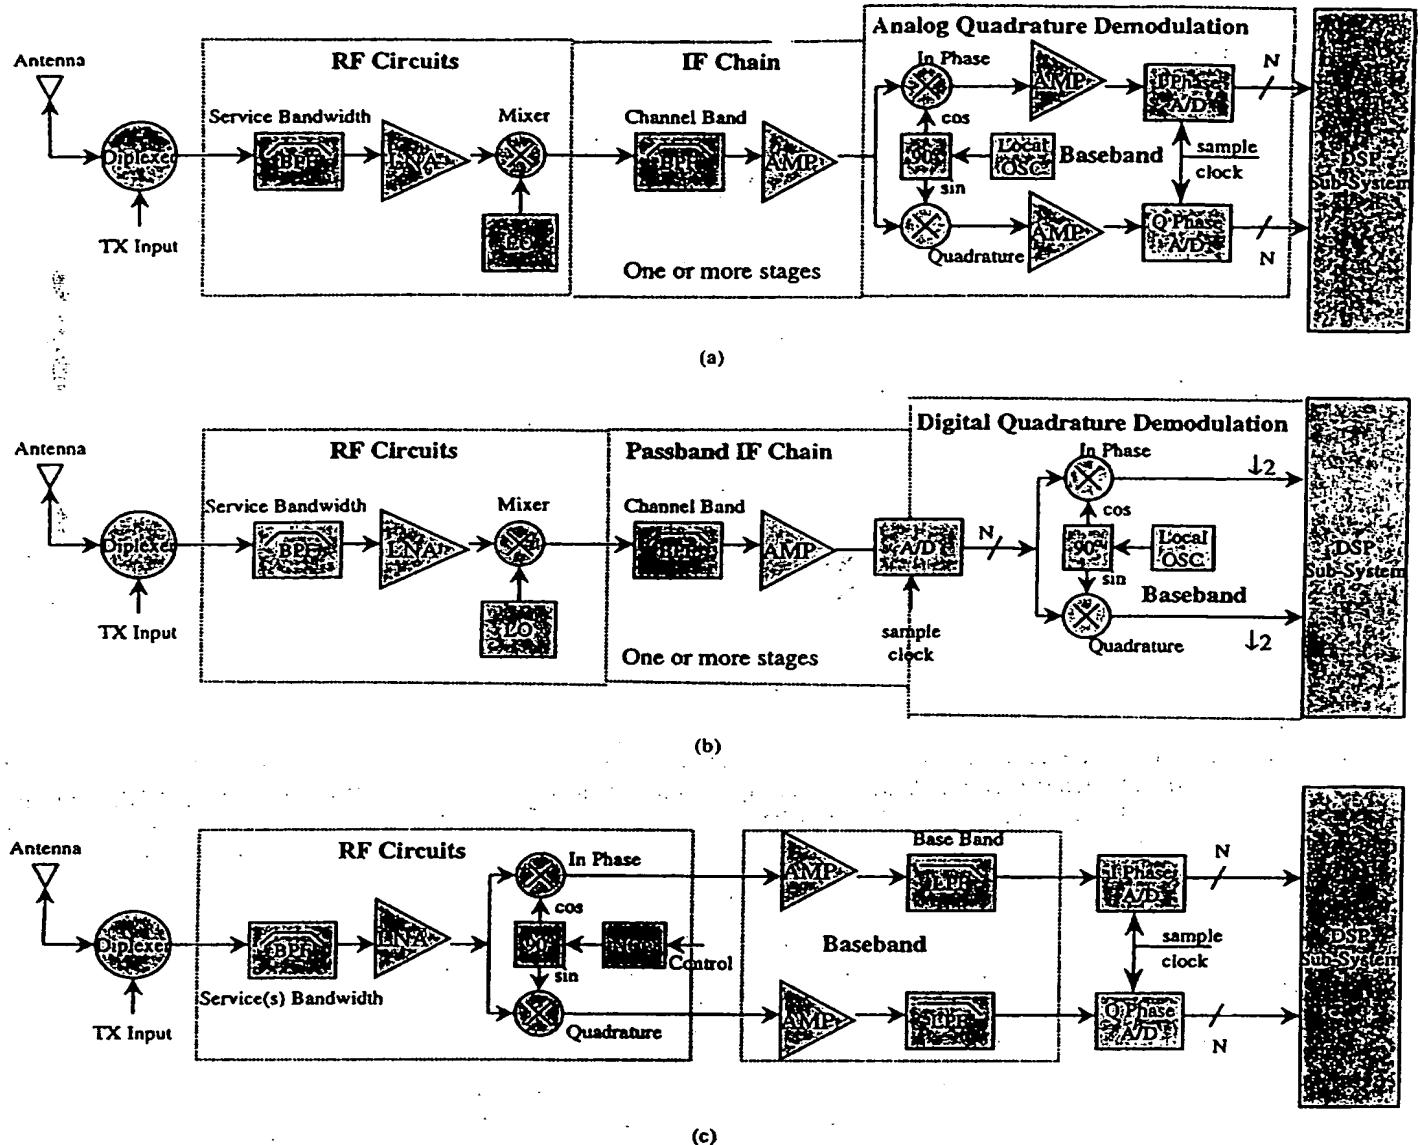


Fig. 4. Common RF receiver configurations for software radios. (a) I, Q superheterodyne (with passband alternatives), (b) passband superheterodyne, and (c) direct conversion (zero IF or homodyne).

least significant bit is greater than ambient noise, sample-and-hold circuits dissipate significant power. Fig. 5 shows the log-log relationship between increased resolution and dissipated power.

It is possible to reduce ADC power by interleaving the sample-and-hold and quantizer circuits. Interleaved ADC's repeat functional blocks of the serial ADC [Fig. 6(a)] to distribute functions over parallel sample-and-hold circuits (Fig. 6(b), left) or over parallel quantizers (Fig. 6(b), right). The interleaved circuits can be integrated on monolithic chips. In many cases, greater performance is achieved at lower total power due to the lower circuit frequencies per parallel path. This is achieved at the expense of larger chip area. Parallel circuits may not be the best low-power option for passband ADC's in which the highest IF frequency is much greater than the signal

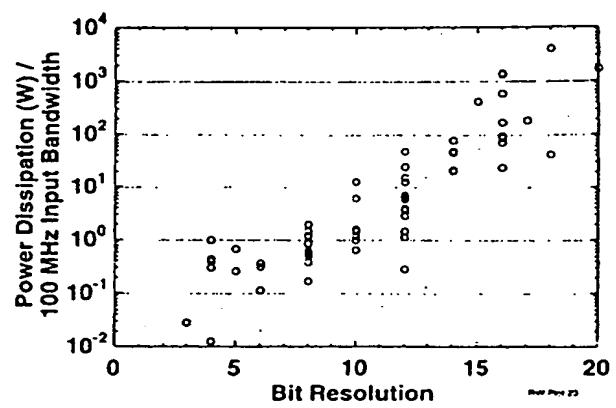


Fig. 5. ADC power dissipation versus resolution [13], [21].

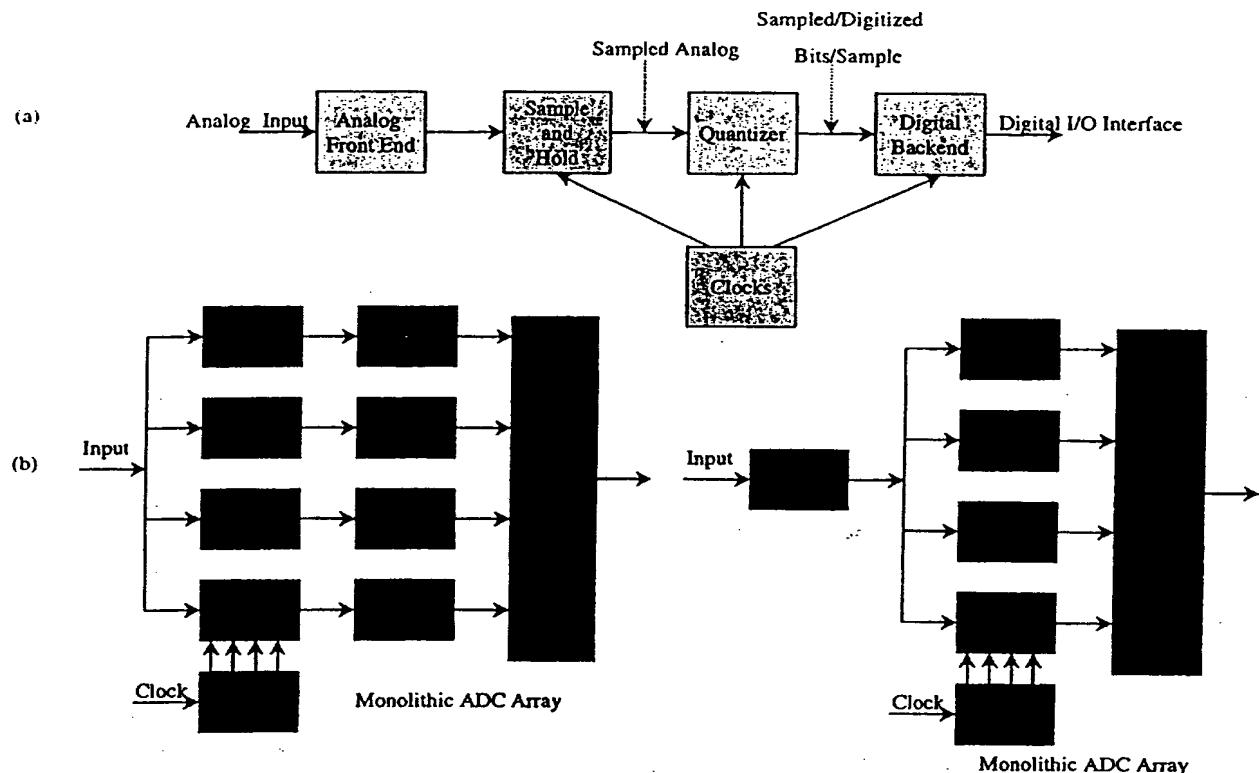


Fig. 6. ADC functional block diagrams. (a) ADC functional diagram and (b) interleaved ADCIS.

bandwidth. In this case, a single sample and hold with parallel quantization yields lower power.

Sample-and-hold circuits are typically implemented in gallium arsenide (GaAs), bi complementary metal-oxide semiconductor (BiCMOS), indium phosphide (InP), hybrid bipolar transistor (HBT), heterojunction field-effect transistor (HFET), and complementary metal-oxide semiconductor (CMOS). GaAs and InP provide high performance with high power dissipation. HBT and HFET circuits have the highest upper cutoff frequencies reported in the literature, and therefore may yield the highest performance ADC's ultimately, but the technology is not yet mature. CMOS traditionally has a lower performance (measured as the product of sampling frequency times the number of bits) than the others, but it dissipates the least power for a given combination of sampling rate and quantization accuracy.

Another approach to low power is the oversampling delta-sigma (also called sigma-delta) converter [14] shown in Fig. 7. The sample-and-hold circuit oversamples the input analog signal by a factor of N . This allows the use of a simpler, lower power converter such as a threshold comparator (a 1 bit converter). In addition, the quality of the antialiasing analog filter (measured as the ratio of -3 dB bandwidth to -23 dB bandwidth) need not be as high as the Nyquist ADC. The cascade of the decimator and digital filter with the antialiasing filter yields a product filter, each stage of which need not be as effective as the single antialiasing filter of the Nyquist ADC.

The delta-sigma and Nyquist converters differ in dissipated power, number of components, and pace of development of the

underlying technology. Since digital technology traditionally advances faster than RF and analog technology in terms of integration and power reduction, delta-sigma converters have a technology advantage. Delta-sigma converters with 2–4 bits of resolution are now deployed in commercial low-power, spread-spectrum wireless terminals. However, increased dynamic range is required for multiple-user interference mitigation, near/far performance, and jammer suppression [15].

In analyzing options for the low-power applications such as SUO, a reasonable aggressive compromise envisions 8–12 bits of dynamic range, 20–30 MHz sampling rate, and 150–300 mW of power dissipation. By the year 2000–2001, the sampling rate is likely to advance to 50–60 MHz, with the same accuracy and dissipated power. It is anticipated that the sample-and-hold circuits would be implemented in CMOS, InP, or HFET with CMOS digital circuits.

V. LOW-POWER DSP DESIGN TECHNIQUES

This section discusses semiconductor power-dissipation trends, sources of CMOS power dissipation, and low-power design techniques.

A. Semiconductor Technology Trends

Fig. 8 shows recent trends in ASIC gate density and power consumption. As the feature size decreases linearly, the number of gates per unit area increases quadratically. As the operating voltage is reduced, power dissipation decreases in proportion to the square of the voltage. At a given operating frequency, the change from 5 to 3.3 V results in a power

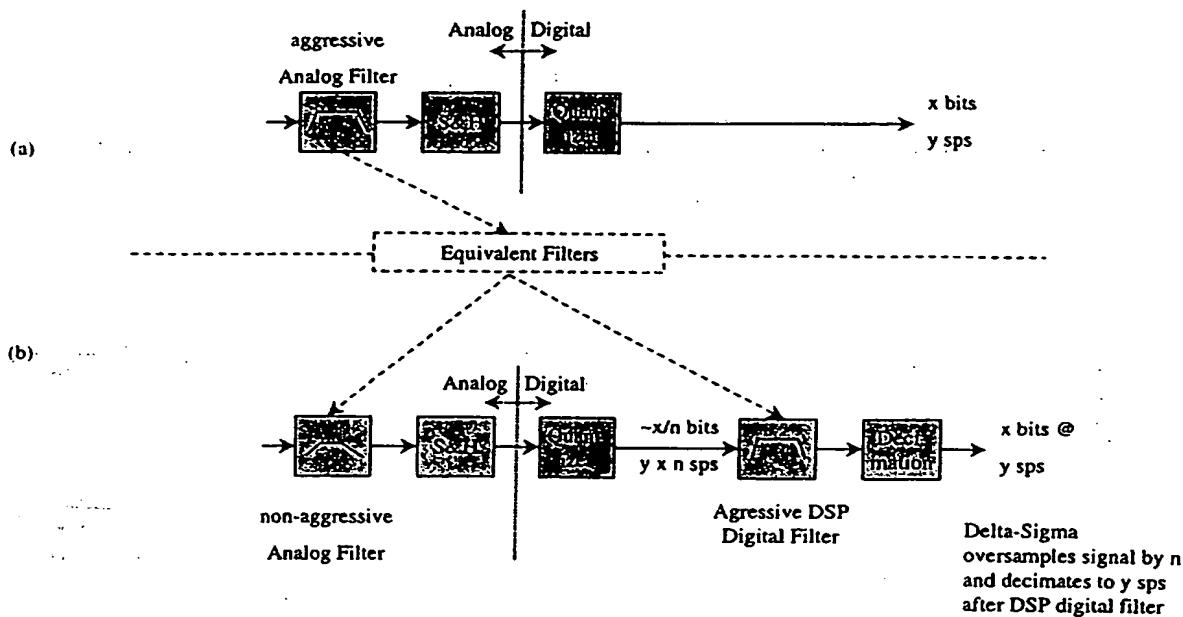


Fig. 7. Nyquist (traditional) versus delta-sigma ADC. (a) Traditional ADC and (b) delta-sigma ADC.

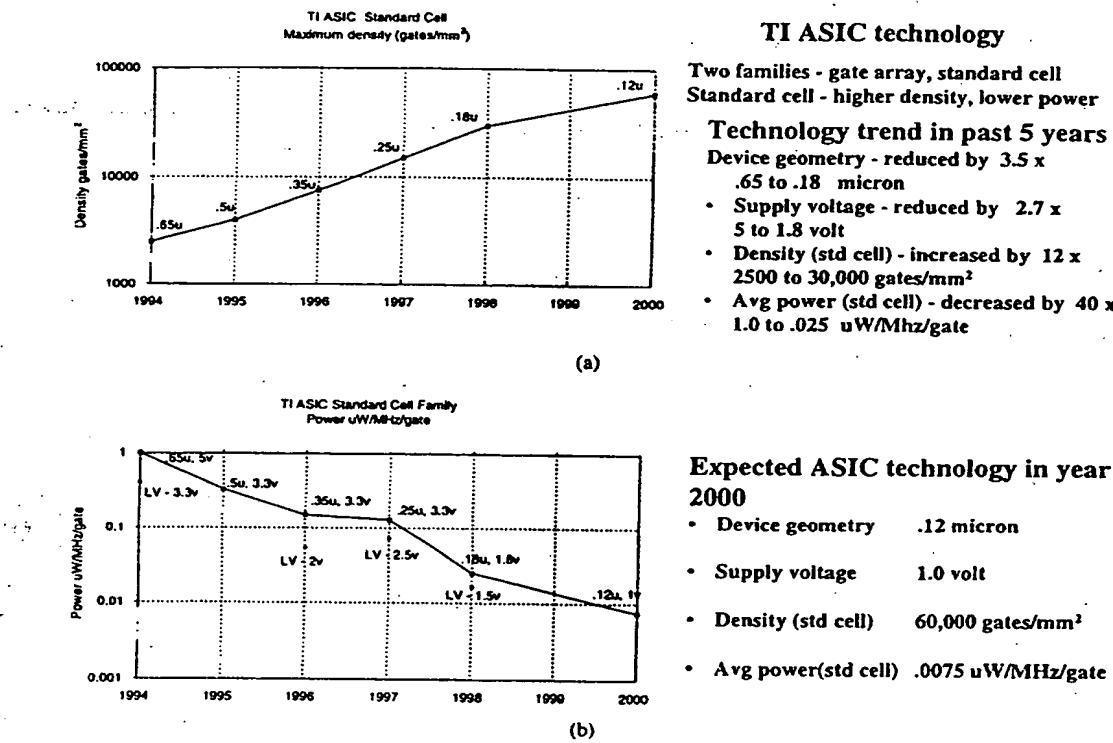


Fig. 8. Texas Instruments ASIC technology trends.

decrease by a factor of 2.5. The change from 3.3 V to lower voltages is already in progress, and will provide an even larger power reduction. A 1.8 V power supply will further reduce power dissipation by a factor of 3.5. And the progression to 1 V technology will reduce power by a factor of about 11 compared with 3.3 V technology.

The standard cell ASIC family shown has higher density and lower power than the gate array family. The most recent

member of the Texas Instruments (TI) ASIC family [5], the TSC6000, has a 0.18- μ m geometry, and operates at 1.8 V. Its density is 30 000 gates/mm², with an average power dissipation of 0.025 μ W/MHz/gate. In comparison to the 5 V devices of about five years ago, the gate density has increased by a factor of 12. The average power per gate has decreased by a factor of 40. This reduction is considerably more than the difference resulting from the square of the supply voltage.

Projecting the expected power dissipation of 1.0 V technology to the year 2000 as the ratio of the square of the supply voltage (1.8–1.0 V) yields an average power dissipation per gate of $0.0075 \mu\text{W}/\text{MHz/gate}$. Density increases to 60 000 gates/mm². Low-power designs of several million gates will therefore be feasible.

B. CMOS Power Dissipation

The power consumption of CMOS circuits includes switching power, short-circuit power, static power, and leakage power. Switching power, the dominant component, is essentially proportional to the square of the supply voltage, and is linearly proportional to the load capacitance, the frequency, and the percentage of time the circuit is active. Short-circuit power depends on the current flowing during switching transients. An increase in transition speed reduces this effect since short-circuit current spikes are present only when multiple circuits are in transition at the same time. Static power is dissipated only when bias currents are present in analog circuits. Leakage power is static power that is unintentionally dissipated by the current in the device during the "off" state. This can become a significant factor in low-voltage designs using low threshold voltages [16]. The low-power ASIC designs for core-based DSP must minimize the effects of the entire range of sources of CMOS power dissipation.

C. Low-Power Design Techniques

A number of techniques can be used to reduce power. Lowering the supply voltage dramatically reduces power, but also degrades performance. Performance can then be enhanced by reducing the threshold voltage. But this results in greater leakage current. Parallel circuits yield higher performance at the expense of greater complexity and chip area. Spurious transitions occur from switching a circuit several times during the same clock cycle because of multiple input changes or differences in signal path length. Power dissipated in spurious transitions can be reduced by latches and by balancing logic paths. Switching can also be postponed by reordering inputs to introduce the most frequently changing inputs later in the logic path. Placement and routing can minimize the product of interconnection capacitance times switching activity through the localization of high-activity networks.

Static random access memory (SRAM) designs that only activate a small portion of the memory array and that use latch-style sense amplifiers will essentially eliminate static current. Power-supply switching transistors that turn off internal power to circuits not in use reduce standby power. For example, high-threshold transistors can turn off power to circuits using low-threshold transistors.

Power management techniques, such as "sleep" and "standby" modes, minimize power during times when only a portion of the circuit is needed. Other portions can then be selectively turned on as necessary, and active circuits can be selected in software or hardware. For example, clock gating can turn off the clock to inactive circuits, based on a sleep bit set by software. Alternatively, clock gating could be based on

an instruction decode that clocks only those circuits required for the instruction.

VI. CONFIGURABLE DIGITAL PREPROCESSORS

The processing capacity needed for wide-band, high-bit-rate, and spread-spectrum waveforms exceeds the capabilities of programmable DSP technology projected during the next five years. Nevertheless, the flexibility needed for software radio implementations can be achieved using configurable digital ASIC's with a microprocessor or DSP that is tailored for software radio applications.

A. The "Ideal" Software Radio Concept

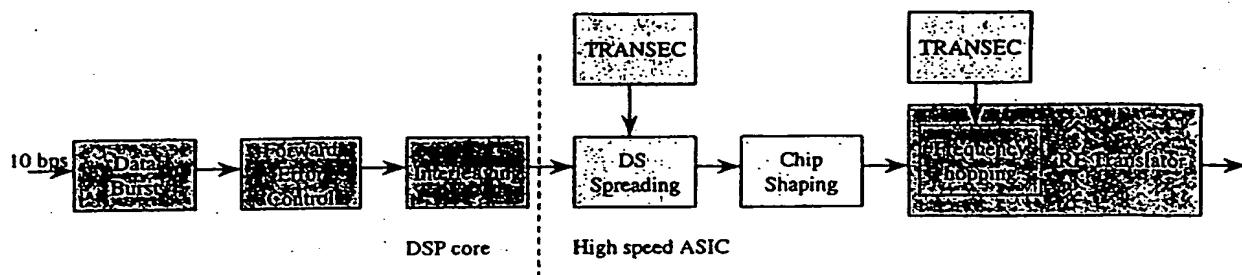
In a "ideal" software radio, practically all of the functions shown in Fig. 1 would be implemented in a general-purpose processor. Transmitter functions ranging from the source encoder to the upconversion of the baseband signal to the final carrier frequency would be performed by this processor. Likewise, the converse functions in the receiver would also be accomplished by the processor, including carrier phase recovery and symbol or pseudonoise (PN) code timing recovery in a spread-spectrum application. In principle, this would allow the same hardware platform to support *any* physical layer imaginable as well as the higher layers of the protocol stack. This ideal radio would only be limited by the capabilities of the analog components (e.g., ADC's, DAC's, power amplifiers, low-noise amplifiers, antialiasing filters, and antenna subsystems), and by the capacity of the processor.

B. The DSP Core and ASIC Approach

The flexibility goal of the software radio can be approximated by:

- moving the ADC and DAC as close to the antenna as possible;
- implementing functions with very high processing demand in ASIC's that can be run-time configured to support a wide range of signal structures;
- maximizing the number of functions performed by the DSP.

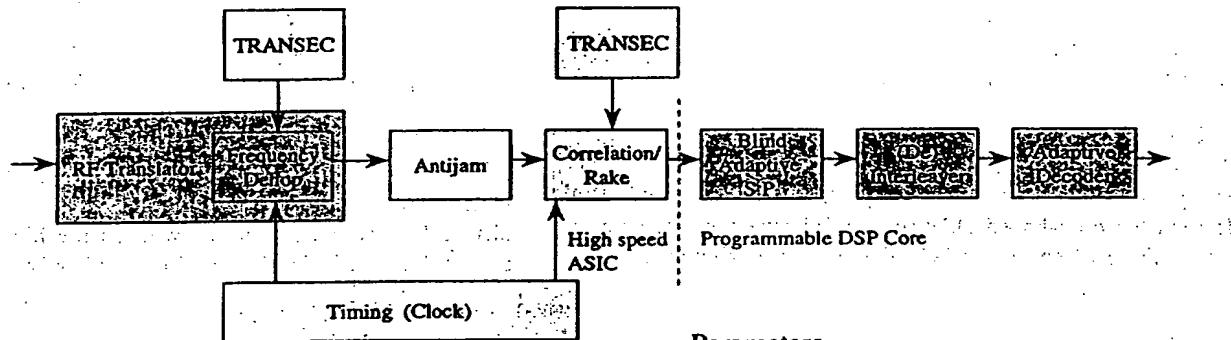
Current DSP, ASIC, and semiconductor technology is rapidly evolving to support DSP core macro cells via an ASIC library [5]. The DSP core is a microprocessor-like programmable DSP that has been designed for efficient gate count, die area, and performance. The DSP core can be augmented with a customized ASIC on a single chip. For software radio applications, the ASIC includes high-throughput modulation functions such as correlators and high-speed filters. It also includes power management and input/output functions. The combination of DSP core and other ASIC functions can be closely tailored to the application. This offers greater potential throughput, lower gate count, smaller die area, and lower power consumption than other semicustom approaches, such as field-programmable gate arrays (FPGA's). Of course, the nonrecurring engineering of the ASIC will be greater than that of an equivalent design using FPGA's and standalone DSP chips.



Parameters

- Burst Rate
- Forward Error Control (Code Parameters, Code Rate)
- Chip Rate
- Hop Rate (# bits/hop)

Fig. 9. DSSS transmitter block diagram. [17].



Parameters

- Burst Rate
- Forward Error Control (Code Parameters, Code Rate)
- Chip Rate
- Hop Rate (# bits/hop)
- Frequency Hop Synthesizer
- Rake (# Fingers)
- Blind Adaptive Signal Processing and/or Pilot Symbols

Fig. 10. DSSS receiver algorithm block diagram. [4], [6].

To some, it seems inappropriate to use the term "flexible" when describing an ASIC because the logic of an *application-specific* chip cannot be modified once the part has been fabricated. However, the generic functions shown in Fig. 1 are common to a broad spectrum of wireless waveforms. Thus, one may be able to design an ASIC that accommodates many waveforms.

A waveform may be characterized by a parameter set defining modulation type, symbol rate, chip rate, pulse shape, constellation, etc. If implemented in an ASIC with programmable parameters, the ASIC can support a variety of waveforms. Consider the pulse-shaping function of the modulate block shown in Fig. 9. Suppose this function is implemented in an ASIC using a finite impulse response filter. If the filter coefficients are stored in RAM, and if this RAM is accessible to the DSP, the DSP can change the pulse shape by changing

the contents of the RAM. A reasonable ASIC design would also allow the number of taps in the filter to be programmable. As long as the number of taps required does not exceed the maximum number available on the ASIC, an arbitrary pulse shape can be supported.

Similar analysis applies to the receiver ASIC of Fig. 10. Timing recovery, transmission security (TRANSEC), and a multifingered correlation receiver that is programmable probably consumes more chip area and power than its waveform-specific counterpart. It is therefore clear that significant effort will be expended in the design of such ASIC's. The implementation of parameterized ASIC's might not be as efficient with respect to gate count and power consumption as that of a point design. Thus, waveform flexibility requires tradeoffs that include waveform types supported, gate count, die area, power consumption, and cost.

C. Waveform Supportability

For handheld radios, size, power, and weight constraints generally limit the amount of processing capacity available for hosting the functions of Fig. 1 in software. This implies that the computationally intensive functions must be implemented in ASIC's. A partitioning of DSSS/CDMA waveforms onto ASIC and DSP components is presented in Figs. 9 and 10. The most crucial parameters in determining waveform supportability by the DSP subsystem include:

- ADC sampling rate;
- dynamic range;
- processing capacity requirement for the translation of digital IF signals to baseband;
- processing capacity requirements of modulation and demodulation algorithms;
- processing capacity requirements of error coding and decoding algorithms (especially, e.g., Viterbi decoders);
- processing capacity of synchronization algorithms—especially the demanding burst mode and satellite applications.

Partitioning a function for ASIC's generally emphasizes increased throughput, but can also emphasize lower power consumption. The precision of arithmetic operations can be customized through the analysis of dynamic range requirements. Lower precision arithmetic can greatly reduce power consumption in the DSSS correlator, for example.

In addition to the physical layer considerations, protocols can limit the flexibility of the proposed software radio. In particular, it may be difficult or impossible to support burst protocols that place arduous requirements on carrier-phase and symbol (or PN code) timing recovery. Such protocols may place a higher processing demand on the system than can be provided in a DSP architecture. Satellite signal tracking has similar requirements. These timing constraints may be met using dedicated ASIC tracking circuits. Thus, a software radio may support the desired range of waveforms and protocols if RF conversion and digital ASIC's with the requisite flexibility are teamed with an appropriate DSP.

D. A DSP Core-Based Software Radio ASIC

To illustrate these concepts, an archetypal software radio is described. The design consists of a single ASIC that includes wireless modem functions, advanced power management, a Texas Instruments TMS320C6xx programmable DSP core [5], [16], and input/output interface logic. This ASIC (Fig. 11) may be called the multipurpose modem chip (MMC).

The spread-spectrum modem logic within the MMC may be programmed by the DSP core for a wide range of functions, including:

- narrow-band jammer suppression;
- digital-to-digital (D/D) conversion;
- demodulation;
- modulation;
- transmit power management;
- military transmission security (TRANSEC) features;
- power management.

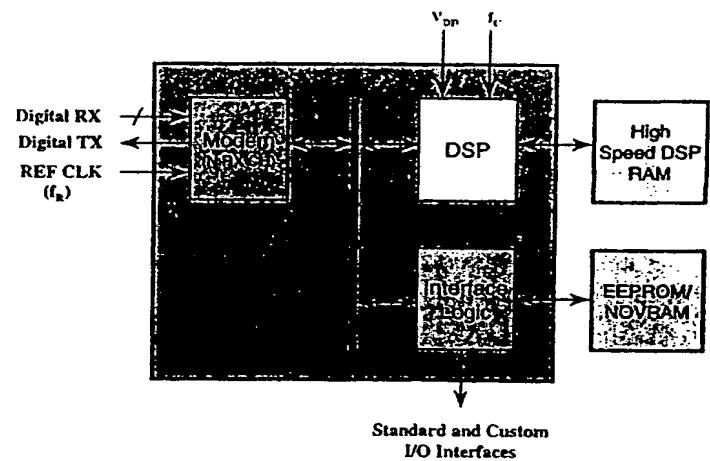


Fig. 11. Multipurpose modem chip [4], [6].

The MMC also generates and demodulates conventional, nonspread-spectrum signals by turning off the PN code generator and mixers. Figs. 9 and 10 show the allocation of the above functions between the ASIC and DSP core. Included in each figure is a list of key parameters that determine whether functions are implemented in the core or elsewhere on the ASIC. With the exception of the military TRANSEC features, the above MMC functions are now described in more detail.

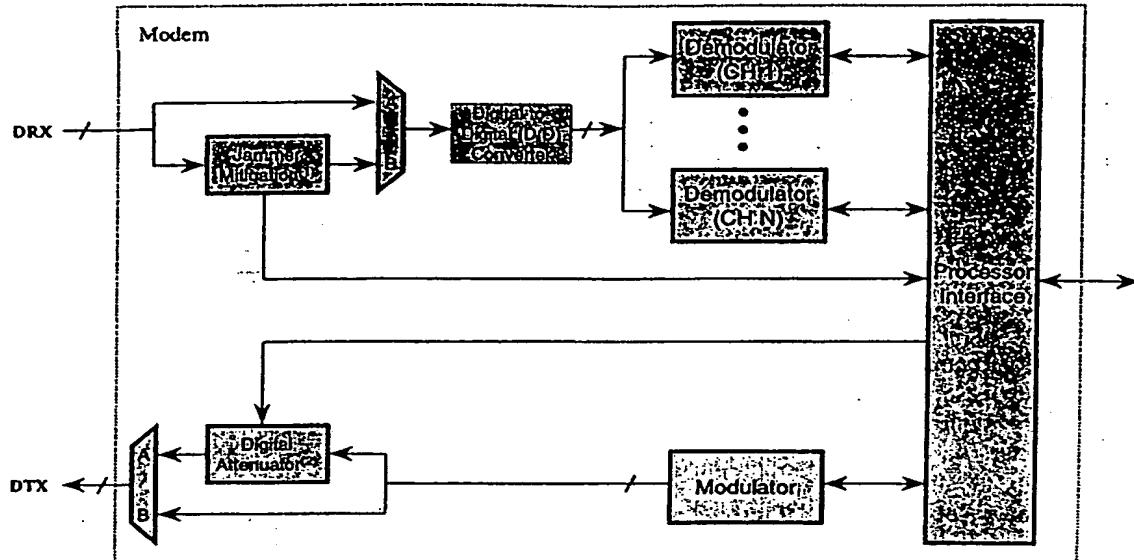
1) Jammer Suppression: Fig. 12 illustrates the main functions of the digital front end of the ASIC including the placement of the jammer suppression block. Mitigation algorithms for this part of the ASIC [15] are chosen to address the characteristics of expected jammers and interference.

The adaptive transversal filter (ATF), for example, excises narrow-band noise and continuous-wave (CW) jammers. The processing for this filter is described in [17] and illustrated in Fig. 13. A typical two-sided ATF design for the MMC calls for 12 bits of precision, 33 taps, and the Widrow-Hoff least mean-square (LMS) algorithm [18], [19] to determine the weight values. The tap weights are maintained in 16 read-only registers within the ASIC, and are accessible to the DSP core. Since the weights are accessible to the DSP core, it computes the ATF transfer function, jammer frequency, and bandwidth in parallel while the ATF is operating.

As shown by the multiplexer (MUX) block after the ATF in Fig. 12, the filter can be bypassed when there is no need for jammer mitigation. Under these circumstances, the low-power ASIC design allows the ATF reference clock to be gated off, thereby conserving power.

2) Digital-to-Digital Conversion: The digital-to-digital converter reduces the number of bits of precision from the jammer mitigation block to the minimum required to demodulate the signal. This substantially reduces the overall gate count and power consumption required for the computationally intensive front-end filtering and correlation functions. Existing DSSS receivers have used lower precision since the dynamic range can be recovered through integration in despreaders.

3) The General Demodulator: The number of independent demodulator channels is determined by the application. In terrestrial systems, at least four channels are needed to im-



- (1) Jammer mitigation includes space-time adaptive antenna processing, an adaptive transversal filter or a combination
- (2) Digital attenuator required for precise transmit power management

Fig. 12. Front-end processing [4], [6].

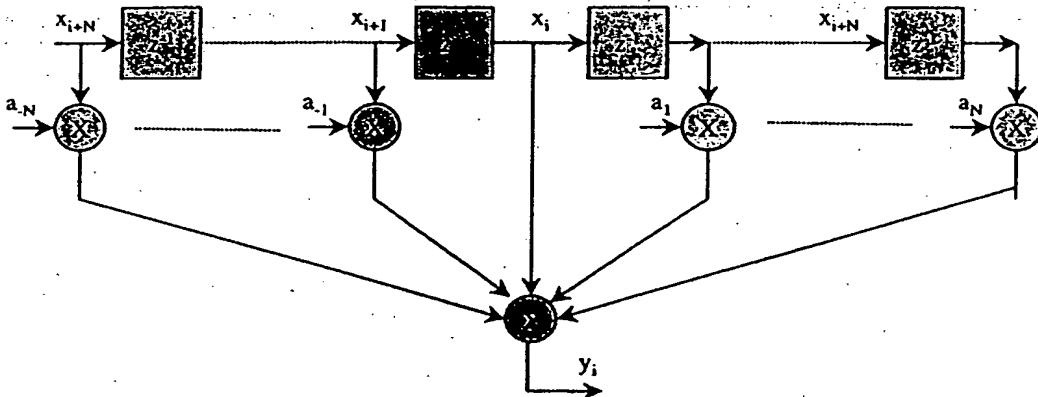


Fig. 13. Adaptive transversal filter.

plement a traditional rake receiver because of nature of the channel. Reception of global positioning satellite (GPS) signals requires as few as four and as many as 12 independent channels. Depending on the number of active channels required, power management provisions in the design allow the DSP to disable the reference clock to unused channels to conserve power.

A more detailed description of the programmable general demodulator is shown in Fig. 14. Conventional digital modulation formats supported by the general demodulator in the MMC design include M -ary phase-shift keying (M-PSK), offset quaternary phase-shift keying (OQPSK), M -ary pulse amplitude modulation (M-PAM), and M -ary quadrature amplitude modulation (M-QAM). Spectral shaping may be used with any of these waveforms. Trellis-coded modulation

(TCM) may be generated and demodulated in the DSP core. Further design tradeoffs include the possibility of a hardware Viterbi decoder in the ASIC. Minimum shift keying (MSK) and Gaussian MSK may also be demodulated primarily using the DSP core.

The first operation performed by each demodulator block of the general demodulator is final quadrature downconversion of the passband signal to baseband. Doppler and reference-oscillator drift are compensated by a numerically controlled oscillator (NCO) having a mean phase and frequency precision of 2^{-16} degrees and 2^{-32} Hz, respectively. The phase and frequency of the NCO are set by the DSP via a memory-mapped I/O (MMIO) interface. The complex baseband samples are then decimated so that matched filter processing uses the lowest possible frequency. The decimator includes a complex

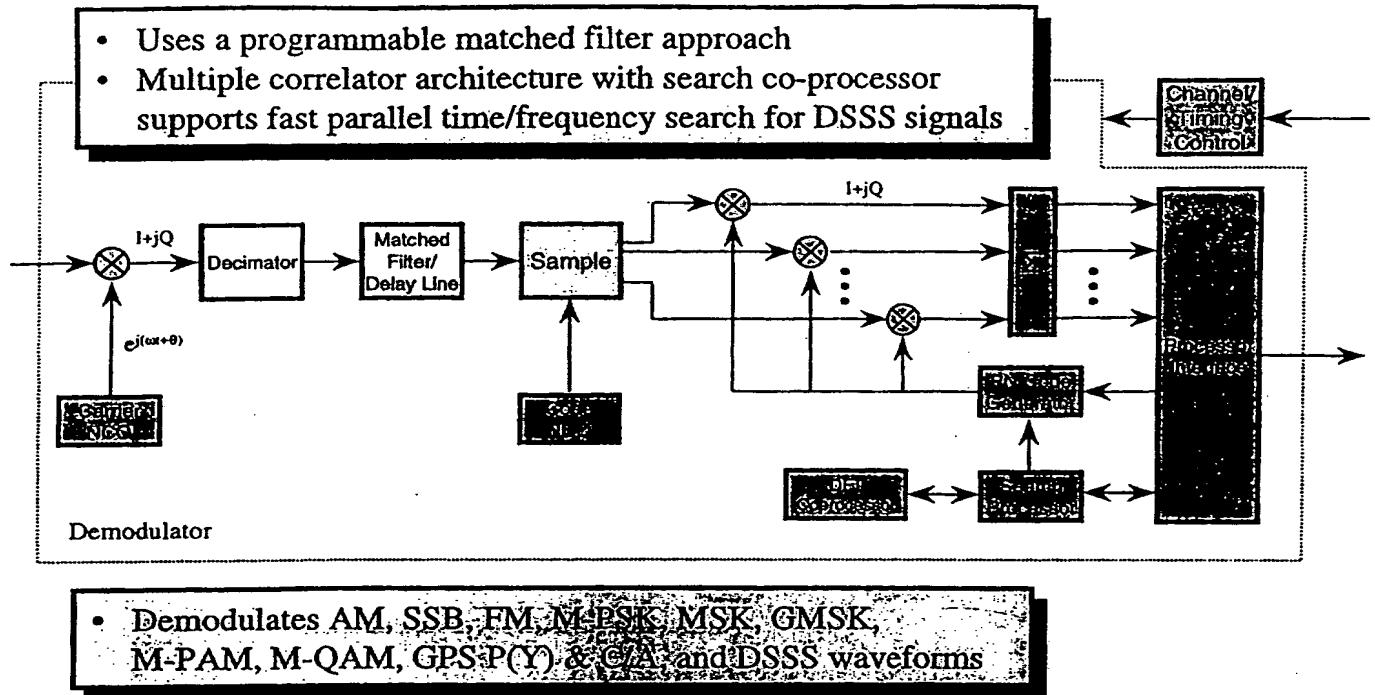


Fig. 14. General demodulator architecture [4], [6].

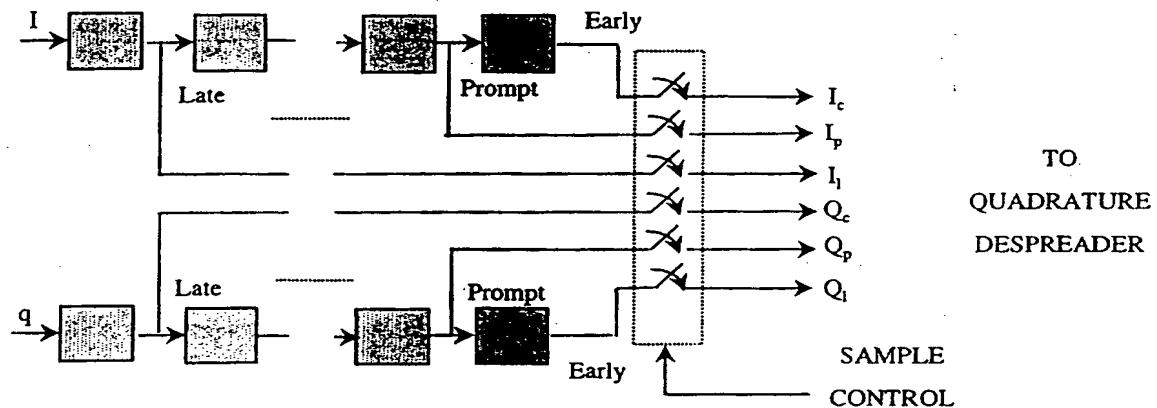


Fig. 15. Delay line and sample functions.

finite impulse response (FIR) filter followed by a decimator. The filter coefficients and decimation rate are controlled by the DSP core.

Following decimation, the complex signal is applied to a matched filter, the coefficients of which are specified by the DSP via the MMIO. The output of the matched filter is sent to the delay line and sample functions, which are shown in greater detail in Fig. 15. The PN chip separation associated with each output in the delay line is a function of the final sample frequency at the output of the decimator. In nonspread applications, the delay line can be used to recover symbol timing. The output of the code NCO determines when the switches in Fig. 15 are closed. The outputs of the delay line are then mixed with the PN code generator output and

accumulated during one symbol period or less. At least three code mixers are needed to track the PN code phase using a delay-locked loop (DLL). As many as 32 code mixers could be used in an alternative ASIC design to reduce PN code acquisition time.

The demodulator design allows the DSP to specify precisely when the accumulation period begins, as well as how long the period will be. This flexibility allows the DSP to track the phase of the data edges in satellite communications systems in which propagation path distances are changing rapidly. It can also track Doppler that measurably lengthens and shortens the symbol period, particularly at low data rates. In nonspread-spectrum applications, the PN code is inhibited, allowing the same demodulator structure to accommodate

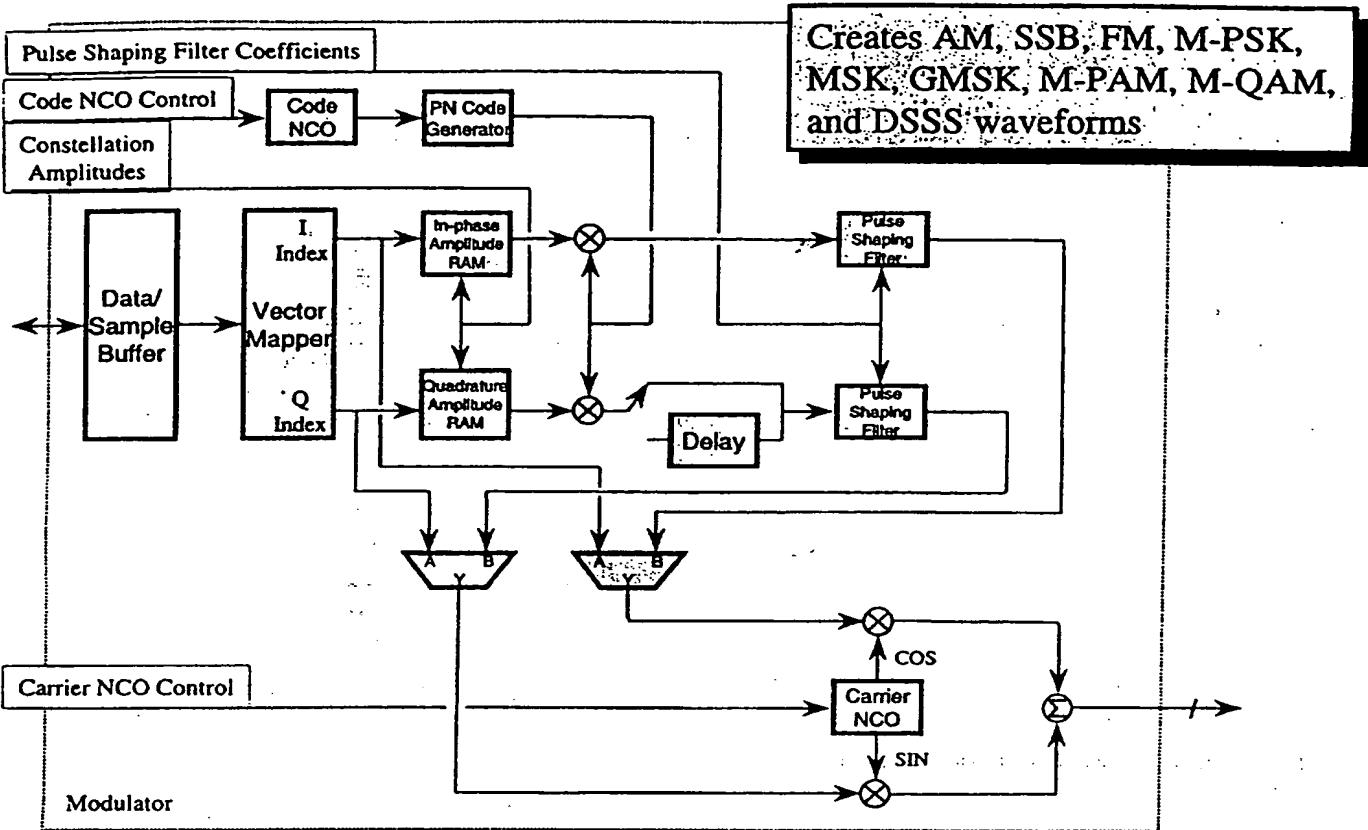


Fig. 16. Programmable general modulator architecture [4], [6].

other waveforms. After accumulation, the I and Q samples are stored in ASIC RAM with an interrupt to the DSP. The DSP reads the I and Q samples via the MMIO interface to complete the demodulation process in software. The code mixer and accumulator (or code correlator) design is such that the DSP can disable the reference clock to unused correlators in order to conserve power. To simplify the logic and conserve gates, code correlators are turned off in groups of four or eight.

Finally, the demodulator contains a search processor and discrete Fourier transform (DFT) coprocessor that accelerate the PN code-timing acquisition process. The search processor forms a time-frequency array of data by computing the DFT of the I and Q samples from each correlator. A sequential probability ratio test is then performed on the time-frequency data to rapidly determine the location of the signal in both time and frequency independently of the DSP core. To initiate a search, the DSP specifies the number of PN chips to be searched and the sample integration period, the inverse of which yields the desired DFT frequency coverage. Then, the search processor signals the DSP with a flag indicating whether the signal was found, and with the PN chip offset and frequency bin of the signal if appropriate. As with other parts of the chip, the search processor may be powered down by the DSP when not in use.

4) *The General Modulator:* The general modulator function of the MMC is illustrated in Fig. 16. It is designed to

Creates AM, SSB, FM, M-PSK, MSK, GMSK, M-PAM, M-QAM, and DSSS waveforms

generate the following baseband signals: amplitude modulation (AM), single sideband (SSB), frequency modulation (FM), M -PSK, MSK, GMSK, M -PAM, and M -QAM. Others may be programmed using the MMC and DSP core.

The modulation process begins when the DSP writes the coded bits to be transmitted into an ASIC data buffer via the MMIO interface. The vector mapper arranges these bits into an index, which selects the appropriate constellation point from the constellation lookup table. These constellation points have been previously initialized by the DSP core. The I and Q samples associated with the selected constellation point are mixed with the applicable PN code sequence. They are then interpolated to the transmit clock frequency, exciting two pulse-shaping filters. These filter coefficients are also initialized by the DSP. A programmable delay can be inserted into the quadrature signal path as needed to generate offset QPSK, MSK, or GMSK.

The outputs of the pulse-shaping filters are then fed through a multiplexer to a quadrature upconverter consisting of the carrier NCO and two mixers. The quadrature upconverter operates like the quadrature downconverter in the demodulator. If the sample buffer contains complex samples of an analog waveform, the digital waveform logic is bypassed, feeding the vector mapper output directly to the quadrature upconverter. The DSP core programs the frequency and phase of the quadrature upconverter and the PN code. This allows the radio

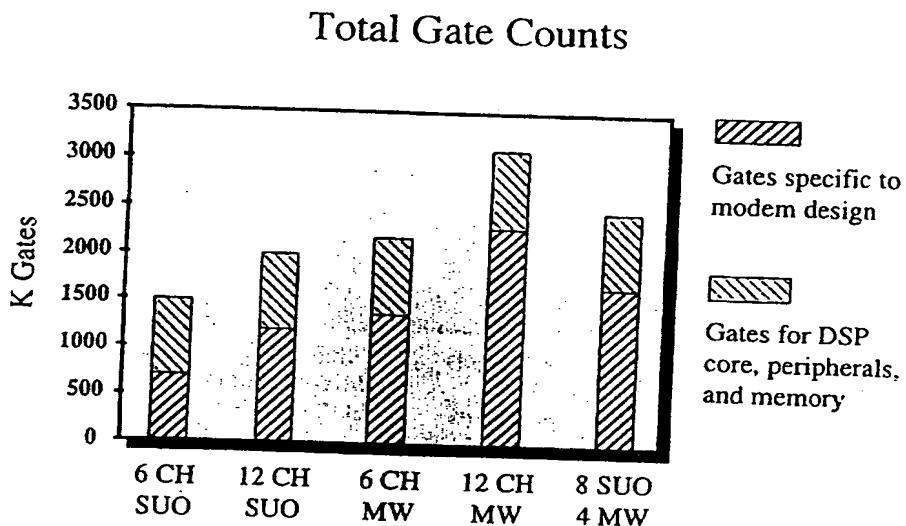


Fig. 17. Total MMC gate count [4].

to perform Doppler precorrection and to correct reference clock drift.

5) *Precision Programmable Attenuator*: The attenuator of Fig. 12 scales the modulator output amplitude using a fixed-point multiplier, the value of which is set by the DSP core. This assures precise power management when required. The numerical precision of the attenuator depends on the dynamic range over which the output power must be controlled. If the dynamic range is large, the precision of the attenuator output could exceed that of the modulator, distorting the waveform. The attenuator block may be bypassed, sending the output of the modulator directly to the DAC.

VII. PARTS COUNT AND POWER CONSUMPTION

A gate-count estimate was generated for a single MMC ASIC incorporating a high-performance DSP core. The receiver complexity drives gate count, which is nearly linear in the number of parallel receiver channels. In addition, the "SUO" configurations include the complex DSSS waveform and a GPS receiver. Multiple-waveform ("MW") configurations accommodate multiple conventional waveforms. Device configurations comparing gate count therefore varied the number of receive channels and the mix of SUO and MW waveforms, as shown in Fig. 17. The corresponding gate counts show the fixed gates required for the DSP core and the variable gates required for the ASIC portions of the MMC design. The number of gates required for 6 MW channels is about the same as the number required for 12 SUO channels. Twelve MW channels requires more than 3 million gates. The mixed configuration of eight SUO channels and four MW channels reduces the gate count to about 2.5 million while providing much of the capability of the 12 MW configuration.

The gate count of the DSP core is based on the C6xxx series high-performance DSP chip [20]. This core has an instruction cache and internal data memory with minimal local memory and DMA interfaces.

The matched filter and increased numerical precision of the correlators increase the gate count of the MW channels. Demodulators were sized with 32 correlators each. For the SUO waveform, the correlators are based on an existing design consisting of separate accumulators for I and Q . This consists of a two-stage accumulator with 7-bit precision in the first stage, followed by a 23-bit accumulator. To save power, the 23-bit accumulator operates at a lower rate, using carries from the first stage accumulator. The precision of the MW first correlator stage was increased to 12 bits.

Fig. 18 shows the estimated power for the demodulator implemented in 1 V technology using a 40 MHz clock. All channels and all 32 correlators per channel are active in the acquisition mode. In the track mode, four channels are used, and only four of the correlators are used per channel.

The six-channel SUO device operates communications and GPS as separate modes, picking one or the other. In each case, all six channels are used for acquisition and four channels are used for tracking. The dissipated power is therefore the same for either GPS or communication mode. The 12-channel MW device may operate the GPS and communication modes simultaneously. The eight SUO channel device could use eight channels for GPS acquisition and four for tracking. At the same time, the four MW channels could be used for communications. In this case, four channels are used for both acquisition and tracking. When all 12 channels are active for acquisition and eight are active for tracking, the device dissipates the power shown in the figure. For the 12-channel MW device, simultaneous operation is also shown; with six channels each for acquisition and four each for tracking. The 12 channel MW demodulator dissipates over 700 mW in the acquisition mode. Track mode power shown in the figure is reasonably low for all of the configurations since the unused channels and correlators are inactive, consuming essentially zero power. These power estimates are based on scaling the power per gate from 1.8 to 1.0 V. There would be some

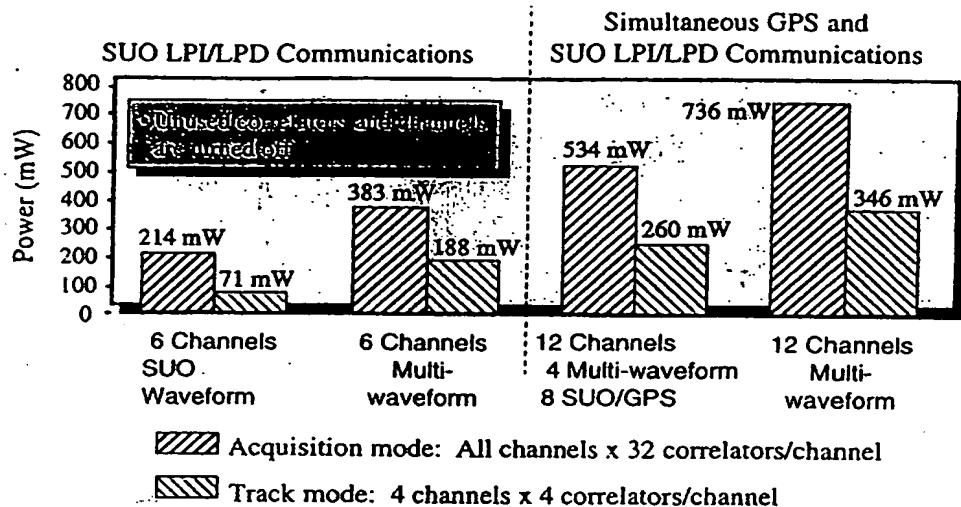


Fig. 18. 1 V demodulator power consumption [4], [6].

- Demodulator, ARM, DSP Core & Peripherals, Memory, and Device I/O

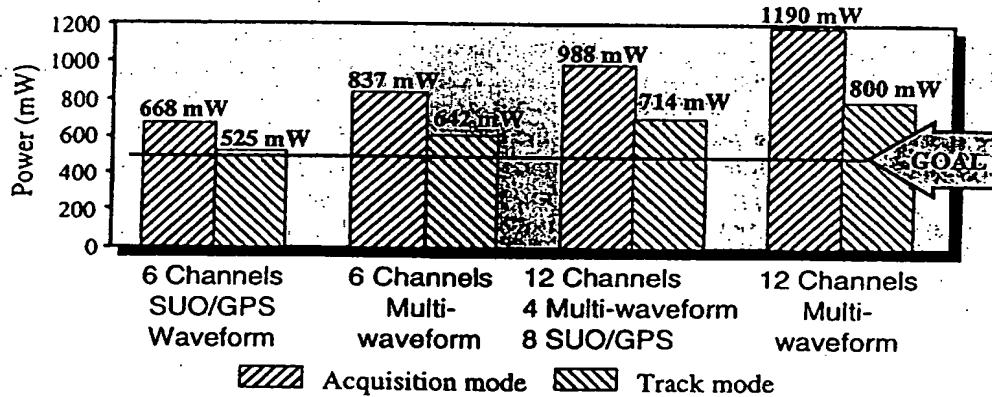


Fig. 19. Receive mode power consumption [4], [6].

additional benefit to reducing feature size from 0.18 to 0.13 or 0.12 μm .

Fig. 19 shows the estimated power for the entire device in the receive mode. The acquisition and track modes are the same as described above. The DSP core operates at 160 MHz and dissipates 440 mW. Its advanced RISC machine (ARM) core operates at 80 MHz to dissipate 12 mW.

VIII. DESIGN METHODOLOGY

The design of low-power DSP-based ASIC's for high-performance applications requires appropriately structured tradeoffs. Advanced wide-band waveforms such as those contemplated for DARPA's SUO program present large processing demand. Satisfying this demand at low power requires attention to power dissipation throughout the design of ASIC's like the MMC. The design methods used to develop this architecture are shown in Fig. 20. Significant

tradeoffs consider algorithms that are optimum from a performance perspective but that dissipate excessive power versus suboptimum algorithms that dissipate less power. Programmable DSP cores yield less throughput or consume more power than configurable ASIC hardware, but at greater cost. In each design, power consumption must be minimized per gate, function, and subsystem, in the context of overall power management strategies such as clock control, sleep, and off modes.

The architecture presented in this paper is motivated mainly by military requirements that include wide-band spread-spectrum waveforms, as well as legacy military waveforms. But the approaches used to develop this architecture are also relevant to commercial applications. However, the specific device designs might not be cost effective for commercial applications. Table III compares military and commercial communication goals. Military requirements

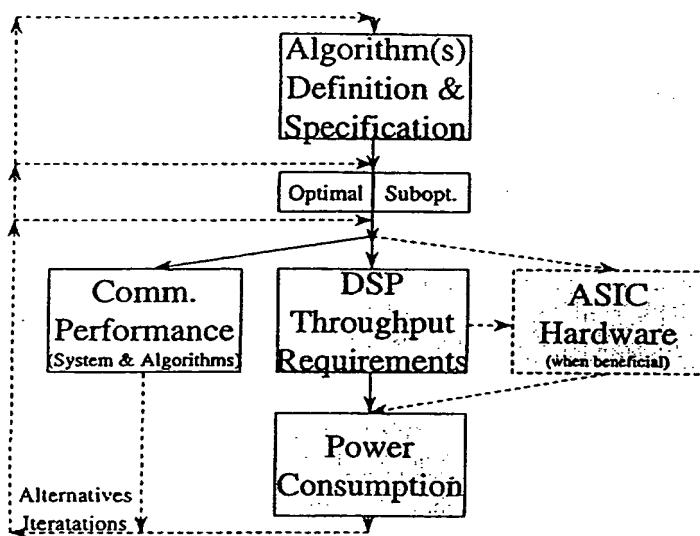


Fig. 20. Low power ASIC development methodology [4], [6].

TABLE III
COMPARISON OF MILITARY AND COMMERCIAL COMMUNICATION GOALS

MILITARY	COMMERCIAL
Peer-to-peer (no base station) operation mode	Basestation/mobile operation/control
Basestation/mobile operation/control	Fixed network infrastructure
Low probability of intercept	Maximum channel capacity
Low probability of detect	Multi-user interference mitigation
Special PN code requirements	Pilot channel for synchronization
Jammer suppression	

for security force hardware commitment to PN codes and channel modulations with features such as low probability of intercept and detection that may be inappropriate to commercial applications. Jammer suppression may also be less appropriate for the commercial sector. In addition, this paper has not addressed emerging multicarrier [22] and wavelet-based [23] communication waveforms. The architecture can support such emerging waveforms provided necessary ASIC circuits are provided. Fast Fourier transform (FFT) circuits may be necessary for multicarrier waveforms. Wavelet waveforms may require filter bank hardware.

IX. CONCLUSION

Commercial and military spread-spectrum waveforms require DSP core and ASIC architectures incorporating advanced power-management techniques. This paper has provided an overview of DSP core-based ASIC designs using the MMC as an example. The result is a complex system-on-a-chip component that implements many of the principles of the software radio in an ASIC that may be implemented with contemporary technology.

REFERENCES

- [1] J. Mitola, "The software radio architecture," *IEEE Commun. Mag.*, vol. 33, pp. 26-38, May 1995.
- [2] B. Sklar, *Digital Communications*. Englewood Cliffs, NJ: Prentice-Hall, 1988, p. 1.
- [3] P. Clark and D. Lammes, "The world phone," *EE Times*, Dec. 8, 1997.
- [4] Adapted from Raytheon TI System (RTIS), Purdue University, and University of Michigan; DARPA Small Unit Operation (SUO) Software Radio and Algorithms Development Program: Open Review 2 Presentation, Nov. 18, 1997.
- [5] Texas Instruments Web site at <http://www.ti.com/sc/docs/news/1996/96037a.htm#support>.
- [6] Adapted from Raytheon TI System (RTIS), Purdue University, and University of Michigan; DARPA Small Unit Operation (SUO) Software Radio and Algorithms Development Program: Open Review 1 Presentation, Sept. 17, 1997.
- [7] A. Papoulis, *The Fourier Integral and Its Applications*. New York: McGraw-Hill, 1962.
- [8] A. Kostic and S. Seetharaman, "Digital signal processors in cellular radio communication," *IEEE Commun. Mag.*, vol. 35, pp. 22-35, Dec. 1997.
- [9] B. Razavi, "Recent advances in RF integrated circuits," *IEEE Commun. Mag.*, vol. 35, pp. 36-43, Dec. 1997.
- [10] G. Siix, "Beam it up: Micromechanical beams may enable a radio on a chip," *Sci. Amer.*, Sept. 1997.
- [11] C. M. Goldsmith, "Performance of low-loss RF-MEMS capacitive switches," *IEEE Microwave Guided Wave Lett.*, vol. 8, Aug. 1998.
- [12] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing*. New York: Macmillan, 1992, ch. 6.
- [13] R. H. Walden, "Analog-to-digital converter survey and analysis," this issue, pp. 539-550.
- [14] J. Candy and G. Ternes, *Oversampling Delta-Sigma Data Converters: Theory, Design, and Simulation*. New York: IEEE Press, 1992.
- [15] L. B. Milstein, "Interference rejection techniques in spread spectrum communication," *Proc. IEEE*, vol. 76, June 1988.
- [16] W. Lee *et al.*, "A 1 V DSP for wireless communication," in *ISSCC Dig. Tech. Papers*, 1997, pp. 92-93.
- [17] L.-M. Li and L. Milstein, "Rejection of narrow-band interference in PN spread spectrum systems using transversal filters," *IEEE Trans. Commun.*, vol. COM-30, May 1982.
- [18] J. G. Proakis, "Adaptive filters," in *Proc. NATO Adv. Study Inst. Commun., Random Process Theory*, 1977, pp. 661-678.
- [19] J. M. McCool and B. Widrow, "Principles and applications of adaptive filters: A tutorial review," in *Proc. IEEE Int. Symp. Circuits Syst.*, 1980, pp. 1143-1157.
- [20] Texas Instruments Web site at <http://www.ti.com/sc/docs/dsp/products/c6000/index.htm>.
- [21] A. C. Seabaugh, Raytheon TI Systems Nanoelectronics Technol. Dept., private communication.
- [22] S. Hara and R. Prasad, "Overview of multicarrier CDMA," *IEEE Commun. Mag.*, pp. 126-133, Dec. 1997.
- [23] A. M. Akansu *et al.*, "Wavelet and subband transforms: Fundamentals and communication applications," *IEEE Commun. Mag.*, pp. 104-115, Dec. 1997.

James E. Gunn (S'63-M'66-SM'97) received the B.S.E.E. and M.S.E.E. degrees from Oklahoma State University, and the Ph.D. degree from Southern Methodist University, specializing in electrical engineering.

He has more than 20 years of industrial experience in communication, telecommunication, signal processing, and control. Functionally, he has contributed as a System Engineer, Software/Firmware Engineer, Technical Marketing Specialist, and Engineering Manager. His most recent activities have focused on advanced software radio architectures for military and commercial waveforms on DARPA's SUO program. Other recent activities have included a DSP subsystem for a software transceiver for the ARDIS wireless network and a low-cost portable DSP-based direction-finding RF receiver. He is coauthor of *Communication Mediums for Intelligent Transportation Systems* (TRB/National Academy Press), which is a multimedia/multimedium communication system design guide.

Kenneth S. Barron received the B.S., M.S., and Ph.D. degrees in electrical engineering from Southern Methodist University, Dallas, TX, in 1984, 1987, and 1995, respectively.

He is currently employed by Raytheon Systems Company, Lewisville, TX. From 1984 to the present, he has worked on various Global Positioning System (GPS) projects. His interests are in the areas of wireless terrestrial and satellite communications systems.

William Ruczcyk received the B.S.E.E. degree from the University of Detroit, Detroit, MI, in 1969.

He has been with Raytheon Systems (formerly the Defense Systems group of TI) since 1969. He has worked in the Special Purpose Radar and Military Computer Departments of TI, where he was involved in the design of moving-target processors and several digital computers. During the past 15 years, he has been involved in the definition and design of a number of ASIC's including a 1750A data processor, Pi-bus interface, a vector processor chip set, and a floating-point DSP. He is a Senior Member of the Technical Staff with the Microelectronics Technology Center, where he is responsible for ASIC architectural definition and design.

THIS PAGE BLANK (USPTO)



Poquet Oliver

Name: Rosa Poquet Oliver
Address: Munich Capitellum, Room C342, Tel: 2911

Userid: RP23103

Printer or e-mail address: P0208

JobNo: PD08006

Format: PS

Batch: 1 of 1

Requested: 07/11/2005 13:53:02

Request Number23817326

Documents

US 5991643A1 I
US 5991643A1TI

Comments

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

BLACK BORDERS

IMAGE CUT OFF AT TOP, BOTTOM OR SIDES

FADED TEXT OR DRAWING

BLURRED OR ILLEGIBLE TEXT OR DRAWING

SKEWED/SLANTED IMAGES

COLOR OR BLACK AND WHITE PHOTOGRAPHS

GRAY SCALE DOCUMENTS

LINES OR MARKS ON ORIGINAL DOCUMENT

REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY

OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)